arXiv:1511.07118v1 [cs.LG] 23 Nov 2015

Cascading Denoising Auto-Encoder as a Deep Directed Generative Model

Dong-Hyun Lee DONGHYUN.LEE.DL@GMAIL.COM
Montreal Institute for Learning Algorithms
Université de Montréal
Montréal, QC, H3C 3J7

Abstract

Recent work (Bengio et al., 2013) has shown how Denoising Auto-Encoders(DAE) become generative models as a density estimator. However, in practice, the framework suffers from a mixing problem in the MCMC sampling process and no direct method to estimate the test log-likelihood.

We consider a directed model with an stochastic identity mapping (simple corruption process) as an inference model and a DAE as a generative model. By cascading these models, we propose Cascading Denoising Auto-Encoders(CDAE) which can generate samples of data distribution from tractable prior distribution under the assumption that probabilistic distribution of corrupted data approaches tractable prior distribution as the level of corruption increases.

This work tries to answer two questions. On the one hand, can deep directed models be successfully trained without intractable posterior inference and difficult optimization of very deep neural networks in inference and generative models? These are unavoidable when recent successful directed model like VAE (Kingma & Welling, 2014) is trained on complex dataset like real images. On the other hand, can DAEs get clean samples of data distribution from heavily corrupted samples which can be considered of tractable prior distribution far from data manifold? so-called global denoising scheme.

Our results show positive responses of these questions and this work can provide fairly simple framework for generative models of very complex dataset.

Proceedings of the 32^{nd} International Conference on Machine Learning, Lille, France, 2015. JMLR: W&CP volume 37. Copyright 2015 by the author(s).

1. Introduction

Denoising Auto-Encoders were originally designed for feature learning algorithm or layer-wise pre-training of deep neural network (Vincent et al., 2010). Recent work (Bengio et al., 2013) has shown that DAEs can also be probabilistic generative models in the context of pure unsupervised learning. They interpret a DAE as a transition operator inducing a Markov chain whose stationary distribution is the data distribution.

However, there are several drawbacks on the framework. At first, sampling process via MCMC chains always suffers from a mixing problem. Starting from 1 data sample, the chain walks through the data manifold, where the walk-step is made up of corruption of a sample and denoising the corrupted sample. If we can de-noise only slightly corrupted samples, that is, with local denoising scheme, the chain can walk only a small step per iteration. This causes a mixing problem in that it is hard to escape from a mode. We can easily guess that generated samples comes from only few modes near the initial sample in spite of good quality of the samples. And sometimes the chain falls in spurious modes if the corruption is too local. (Bengio et al., 2013) suggested the walkback training to avoid it.

Secondly, lack of direct method to estimate test loglikelihood makes it difficult to compare DAEs to other generative models. Though parzen window estimator could be used for any type of generative models, test loglikelihood is still major measure for quality of generative models.

Recently, directed generative models have been the focus of attention, as for example, Variational Auto-Encoder (Kingma & Welling, 2014) which showed great success both as generative models and in the context of semi-supervised learning (Kingma et al., 2014). Directed generative models are appealing because of their clear probabilistic interpretation which allows ancestral sampling and direct methods to estimate test loglikelihood.

So we interpret DAEs as directed generative models with an stochastic identity mapping (corruption process) as an

inference model and a DAE as a generative model in order to alleviate problems mentioned above under the assumption that distribution of corrupted data approaches tractable prior distribution as the level of corruption increases. But it is difficult for single DAE to generate clean samples of data distribution from tractable prior distribution like gaussian because ususally DAE can de-noise slightly corrupted samples near data manifold, local denoising scheme. For global denoising scheme, we should use cascading DAEs in which single DAE only need to make a transition to more probable distribution of data.

Our model can provide much better method to generate samples from DAEs and a direct method to estimate test log-likelihood. Moreover, it can provide new "exhaustive" training method for DAE and it is also helpful for representation learning.

This work tries to answer two questions. On the one hand, can deep directed models be successfully trained without intractable posterior inference and difficult optimization of very deep neural networks in inference and generative models? These are unavoidable when recent successful directed model like VAE (Kingma & Welling, 2014) is trained on complex dataset like real images. On the other hand, can DAEs get clean samples of data distribution from heavily corrupted samples which can be considered of tractable prior distribution far from data manifold? so-called global denoising scheme.

Our results show positive responses of these questions and this work can provide fairly simple framework for generative models of very complex dataset.

2. Method

2.1. Denoising Auto-Encoder as a Directed Generative Model

In general, features obtained from encoder in Auto-Encoder have no probabilistic interpretation without any sampling process on it to match prior distribution. Moreover, encoder and decoder are not distinguished in generalized version of DAE (Bengio et al., 2013). Because sampling process as a corruption is on the observed input variable, we can consider corrupted input $\tilde{\mathbf{x}}$ as a latent variable $\mathbf{z}$ with an inference model of a stochastic identity mapping.

We have the inference model $q(\mathbf{x}, \mathbf{z})$ and the generative model $p(\mathbf{x}, \mathbf{z})$ as following,

$$\begin{aligned} q(\mathbf{x}, \mathbf{z}) &= q(\mathbf{z}|\mathbf{x})q(\mathbf{x}) = c(\mathbf{z} = \tilde{\mathbf{x}}|\mathbf{x})q(\mathbf{x}) \\ p(\mathbf{x}, \mathbf{z}) &= p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z}) \end{aligned} \tag{1}$$

where $q(\mathbf{x})$ is data distribution, $p(\mathbf{z})$ is prior, $c(\mathbf{z} = \tilde{\mathbf{x}}|\mathbf{x})$ is a corruption process, $p_\theta(\mathbf{x}|\mathbf{z})$ is parametric reconstruction mapping in DAE.

To maximize a variational lower bound on the log-likelihood,

$$\begin{aligned} \log p(\mathbf{x}) &= \log \sum_{\mathbf{z}} q(\mathbf{z}|\mathbf{x}) \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \geq \sum_{\mathbf{z}} q(\mathbf{z}|\mathbf{x}) \log \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \\ &= \mathop{\mathbb{E}}_{\mathbf{z} \sim c(\mathbf{z}=\tilde{\mathbf{x}}|\mathbf{x})} \log p_\theta(\mathbf{x}|\mathbf{z}) - \mathcal{D}_{KL}(c(\mathbf{z} = \tilde{\mathbf{x}}|\mathbf{x})||p(\mathbf{z})) \\ &= \mathcal{L}_\theta(\mathbf{x}) \end{aligned} \tag{2}$$

where $\mathbf{z} \sim c(\mathbf{z} = \tilde{\mathbf{x}}|\mathbf{x})$ is a fixed corruption on the same space as the observed variable. Unlike typical variational training, we use a fixed stochastic encoder so that we don't train inference model to approximate the posterior. Instead, we rely on the simple fact that the distribution of corrupted samples become similar to tractable prior distribution like gaussian as the corruption level increases. It can alleviate difficulty of training in that we only need to train reconstruction part, whereas it is challenging to de-noise heavily corrupted sample which is almost same as tractable distribution, so called global denoising scheme.

2.2. Cascading DAE with gradually increasing corruption

Inspired by (Sohl-Dickstein et al., 2015), we introduce Cascading Denoising Auto-Encoders(CDAE). At the first step, we don't need to recover clean data samples at once, but only have to make a transition from prior distribution to more probable distribution of data. Through several steps, we eventually can get clean data samples.

At first, we have corrupted samples with gradually increasing corruption level $\mathbf{x}^{(1)}$, ... $\mathbf{x}^{(k)}$ of data $\mathbf{x}^{(0)}$.

$$\begin{aligned} q(\mathbf{x}^{(0)}, ..., \mathbf{z} = \mathbf{x}^{(k)}) &= q(\mathbf{x}^{(0)}) \prod_{i=1}^{k} q(\mathbf{x}^{(i)}|\mathbf{x}^{(0)}) \\ p(\mathbf{x}^{(0)}, ..., \mathbf{z} = \mathbf{x}^{(k)}) &= p(\mathbf{x}^{(k)}) \prod_{i=1}^{k} p_\theta(\mathbf{x}^{(i-1)}|\mathbf{x}^{(i)}) \end{aligned} \tag{3}$$

where $q(\mathbf{x}^{(0)})$ is a data distribution and $p(\mathbf{z} = \mathbf{x}^{(k)})$ is a prior distribution. To maximize a variational lower bound

on the log-likelihood,

$$\begin{aligned}\log p(\mathbf{x}^{(0)}) &\geq \sum_{\mathbf{x}^{(1\dots k)}} q(\mathbf{x}^{(1\dots k)}|\mathbf{x}^{(0)}) \log \frac{p(\mathbf{x}^{(0)}, \mathbf{x}^{(1\dots k)})}{q(\mathbf{x}^{(1\dots k)}|\mathbf{x}^{(0)})} \\ &= \mathop{\mathbb{E}}_{\mathbf{x}^{(1\dots k)}} \log \left[\frac{p(\mathbf{x}^{(k)}) p(\mathbf{x}^{(0)}|\mathbf{x}^{(1)})}{q(\mathbf{x}^{(k)}|\mathbf{x}^{(0)})} \prod_{i=1}^{k-1} \frac{p_\theta(\mathbf{x}^{(i)}|\mathbf{x}^{(i+1)})}{q(\mathbf{x}^{(i)}|\mathbf{x}^{(0)})} \right] \\ &= \mathcal{L}_\theta(\mathbf{x}) \\ &= \mathop{\mathbb{E}}_{\mathbf{x}^{(1)}} \log p_\theta(\mathbf{x}^{(0)}|\mathbf{x}^{(1)}) \\ &- \sum_{i=1}^{k-1} \mathop{\mathbb{E}}_{\mathbf{x}^{(i+1)}} \mathcal{D}_{KL}(q(\mathbf{x}^{(i)}|\mathbf{x}^{(0)})||p_\theta(\mathbf{x}^{(i)}|\mathbf{x}^{(i+1)})) \\ &- \mathcal{D}_{KL}(q(\mathbf{x}^{(k)}|\mathbf{x}^{(0)})||p(\mathbf{z} = \mathbf{x}^{(k)}))\end{aligned} \tag{4}$$

where all expectations are $\mathbf{x}^{(i)} \sim q(\mathbf{x}^{(i)}|\mathbf{x}^{(0)})$. Different from diffusion process in (Sohl-Dickstein et al., 2015), we use the same DAE in all $p_\theta(\mathbf{x}^{(i)}|\mathbf{x}^{(i+1)})$'s (that is, we use shared weights for all steps like Recurrent Neural Networks), and conditionally independent intermediate corruption distribution given the data sample. These enable us to consider this model as a series of DAE. As you shall see, this objective is merely the sum of denoising reconstruction errors with several corruption levels in the gaussian case.

2.3. Gaussian CDAE for continuous data

For continuous data, isotropic gaussian is the simplest and typical choice as well as it can show how this framework is related with the typical training objective of DAE, mean squared reconstruction error with gaussian corruption.

$$\begin{aligned}p_\theta(\mathbf{x}^{(i)}|\mathbf{x}^{(i+1)}) &= \mathcal{N}(\mathbf{r}_\theta(\mathbf{x}^{(i+1)}), {\sigma_r^{(i)}}^2) \\ q(\mathbf{x}^{(i)}|\mathbf{x}^{(0)}) &= \mathcal{N}(\mathbf{x}^{(0)}, {\sigma_c^{(i)}}^2)\end{aligned} \tag{5}$$

The negative variational lower bound is

$$\begin{aligned}-\mathcal{L}_\theta(\mathbf{x}) = \sum_{i=1}^{k} \mathop{\mathbb{E}}_{\mathbf{x}^{(i)}} \frac{1}{2{\sigma_r^{(i-1)}}^2} \|\mathbf{x}^{(0)} - \mathbf{r}_\theta(\mathbf{x}^{(i)})\|^2 + \phi(\sigma_c^{(i)}, \sigma_r^{(i)}) \\ + \mathcal{D}_{KL}(q(\mathbf{x}^{(k)}|\mathbf{x}^{(0)})||p(\mathbf{z} = \mathbf{x}^{(k)}))\end{aligned} \tag{6}$$

where the first term is merely denoising reconstruction error with several levels of corruption $\mathbf{x}^{(i)} \sim c(\mathbf{x}^{(i)}|\mathbf{x}^{(0)})$, the second term $\phi(\sigma_c^{(i)}, \sigma_r^{(i)})$ is independent from θ so that it disappears the gradient of this objective, the third term is to match between heavily corrupted samples on the top DAE and the prior as a tractable distribution. We can train the parameters of prior (for example, the mean and variance of gaussian) with this term even before training DAE because we use a fixed identity encoder independent from the parameters of DAE. So this is directly connected with typical DAE in that the terms related with the parameter of DAE θ are only reconstruction errors. Moreover, it makes training our model much easier under the assumption that corrupted samples become closer to tractable distribution as the level of corruption increases.

Moreover, the terms of 6 related with θ can shrink even to single reconstruction error.

$$L_\theta^{shrink}(\mathbf{x}) = \mathop{\mathbb{E}}_{\mathbf{x}^{(c)} \sim Q(\mathbf{x}^{(c)}|\mathbf{x}^{(0)})} \|\mathbf{x}^{(0)} - \mathbf{r}_\theta(\mathbf{x}^{(c)})\|^2 \tag{7}$$

where

$$Q(\mathbf{x}^{(c)}|\mathbf{x}^{(0)}) = \sum_i \frac{1}{2{\sigma_r^{(i-1)}}^2} q_i(\mathbf{x}^{(c)}|\mathbf{x}^{(0)}) \tag{8}$$

where $q_i(\mathbf{x}^{(c)}|\mathbf{x}^{(0)})$ is the same distribution as $q(\mathbf{x}^{(i)}|\mathbf{x}^{(0)})$.

Because our DAE is probabilistic models with sampling after reconstruction, we need to determine $\sigma_r^{(i)}$'s. these variables can be trained with θ, instead, we calculate them before every training epoch thanks to following analytic solution of $\partial \mathcal{L}_\theta(\mathbf{x}) / \partial \left({\sigma_r^{(i)}}^{-2} \right) = 0$

$${\sigma_r^{(i)}}^2 = {\sigma_c^{(i)}}^2 + \frac{1}{d} \|\mathbf{x}^{(0)} - \mathbf{r}_\theta(\mathbf{x}^{(i+1)})\|^2 \tag{9}$$

where d is dimensionality of data and $\sigma_c^{(i)}$'s are gradually increasing values as hyperparameters.

2.4. Estimating test loglikelihood

Due to lack of a direct method to estimate test loglikelihood of generative DAE (Bengio et al., 2013), they use parzen window estimator. It is convenient and unique method for generative models in which the exact loglikelihood is not tractable. But the estimator has several drawbacks (....) Our framework provides the direct method to estimate test loglikelihood which is the most popular and historic one to measure quality of generative models.

We can calculate the variational lower bound of loglikelihood using equation 4, The more the number of step is, the larger the gap between the exact value and the bound is. The gap is

$$\mathcal{D}_{KL} \left(q(\mathbf{x}^{(1..k)}|\mathbf{x}^{(0)})||p(\mathbf{x}^{(1..k)}|\mathbf{x}^{(0)}) \right) \tag{10}$$

and this value can be larger as the number of steps k increases.

In RAISE (Burda et al., 2015) and AIS (Neal, 2001), they estimate the exact loglikelihood through importance weights. We also can do it as following,

$$\begin{aligned}\log p(\mathbf{x}^{(0)}_{test}) &= \log \sum_{\mathbf{x}^{(1\dots k)}} q(\mathbf{x}^{(1\dots k)}|\mathbf{x}^{(0)}_{test}) \frac{p(\mathbf{x}^{(0)}_{test}, \mathbf{x}^{(1\dots k)})}{q(\mathbf{x}^{(1\dots k)}|\mathbf{x}^{(0)})} \\ &= \log \mathop{\mathbb{E}}_{\mathbf{x}^{(1\dots k)}} \left[p(\mathbf{z} = \mathbf{x}^{(k)}) \prod_{i=0}^{k-1} \frac{p_\theta(\mathbf{x}^{(i)}|\mathbf{x}^{(i+1)})}{q(\mathbf{x}^{(i+1)}|\mathbf{x}^{(0)}_{test})} \right] \end{aligned} \tag{11}$$

Acknowledgments

References

Bengio, Yoshua, Yao, Li, Alain, Guillaume, and Vincent, Pascal. Generalized denoising auto-encoders as generative models. In *NIPS'2013*, 2013.

Burda, Yuri, Grosse, Roger B., and Salakhutdinov, Ruslan. Accurate and conservative estimates of mrf log-likelihood using reverse annealing. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics. 2015.*, 2015.

Kingma, D.P., Rezende, D.J., Mohamed, S., and Welling, M. Semi-supervised learning with deep generative models. In *NIPS'2014*, 2014.

Kingma, Durk P. and Welling, Max. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.

Neal, Radford M. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, April 2001. URL http://dx.doi.org/10.1023/A:1008923215028.

Sohl-Dickstein, Jascha, Weiss, Eric A., Maheswaranathan, Niru, and Ganguli, Surya. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 31th International Conference on Machine Learning (ICML'15)*, 2015.

Vincent, Pascal, Larochelle, Hugo, Lajoie, Isabelle, Bengio, Yoshua, and Manzagol, Pierre-Antoine. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J. Machine Learning Res.*, 11, 2010.